View-dependent Radiance Caching

Yangyang Zhao* McGill University
Laurent Belcour† Unity Research
Derek Nowrouzezahrai‡ McGill University

ABSTRACT

Radiance caching is used to accelerate global illumination computations, exploiting the spatial coherence of indirect illumination on surfaces. We propose a new radiance caching approach capable of more correctly reconstructing inter-reflections between glossy surfaces, all while improving performance compared to previous approaches. Contrary to previous works, our view-dependent radiance caching scheme does not heavily rely fundamentally on basis-space representations such as spherical harmonics, and can directly treat outgoing radiance at surfaces instead of incoming radiance distributions. We introduce a new view-dependent record placement strategy and adapt recent Hessian-based error metrics to our view-dependent records [8]. To do so, we derive and compute more accurate derivatives of radiance at surfaces in the scene.

Keywords: Ray tracing, Global illumination, Radiance Caching, Spherical harmonics.

Index Terms: Computing methodologies—Computer graphics—Rendering—Ray tracing;

1 INTRODUCTION

Efficiently computing accurate global illumination effects remains an important problem in computer graphics. Computing indirect illumination, caused by secondary bounces of light off of surfaces, is particularly challenging. These costly effects often require more complex path-based methods, such as path tracing [4] or photon mapping [3] in order to solve the rendering equation using stochastic Monte Carlo ray tracing.

Instead of focusing on efficiently computing global illumination at any specific shade point, irradiance caching [10] accelerates rendering by exploiting the spatial correlation between diffuse indirect shading. Radiance caching was later introduced as an extension that treats glossy surfaces [5]. Here, spherical harmonics (SH) are used to approximate the incoming radiance distribution at the cache locations, instead of simply storing scalar diffuse (ir)radiance. At run-time, the incoming radiances SH coefficients are interpolated and used to reconstruct outgoing radiance at shade points.

The manner in which cache records are placed on surfaces is crucial to the performance of irradiance and radiance caching. Several works address the cache placement problem. In the case of irradiance cache record placement, an appropriate maximum separation between records can be computed according to a *Split Sphere* heuristic [10]. Ward and Heckbert later update this heuristic to leverage a conservative estimate of the irradiances translational gradient [9]. More recently, Jarosz et al. [2] and Schwarzhaupt et al. [8] introduce more accurate visibility-aware Hessian-based heuristics for adaptive cache record placement, significantly improving the quality and performance of the previous approaches. Despite these developments in the irradiance caching context, more advanced cache placement strategies for radiance caching have not been well studied: Krivanek [5] use a strategy based on the gradient of the spherical harmonic coefficients and Scherzer et al. [7] place records according to a blue-noise distribution on surfaces.

We extend Hessian-based error metrics from an irradiance caching context to radiance caching, significantly improving record placement and reconstruction. We propose caching the outgoing radiance instead of the SH incident radiance, and we demonstrate how to improve the ability to deal with complex indirect radiance distributions that result from glossy-glossy surface interactions. Our new radiance caching algorithm is robust, accurate, and adapts to the spatial and angular variation in radiance at and between cache records.

2 PREVIOUS WORK

Ward et al. [10] exploit the fact that the indirect illumination on a diffuse surface changes slowly across the surface. They propose *irradiance caching* to accelerate the computation of the indirect illumination component on diffuse surfaces, a method that caches accurate indirect illumination at a small subset of points in the scene and then interpolates between these cache records.

Ward and Heckbert [9] use translational and rotational gradients of the irradiance to improve record placement. They also propose a first-order Taylor expansion of the irradiance between records in order to better interpolate the irradiance from record locations to arbitrary shade points. Their approach was further improved by Krivanek [5], where a new gradient formulation that relies on stratification-based gradients was employed.

The recent work of Jarosz et al. [2] introduced a visibility-aware Hessian-based error metric to improve the placement of irradiance caching records, and Schwarzhaupt et al. [8] propose a triangulation heuristic to treat occlusion changes when accounting for the gradients and Hessians on diffuse surfaces. They demonstrate significant improvements in the quality and performance of irradiance caching, highlighting the importance of treating visibility discontinuities during placement and interpolation. These approaches, however, only apply to irradiance caching and do not effectively transpose to the radiance caching context.

Krivaneks seminal works on radiance caching [5] were motivated with the added complexity that arises when treating indirect illumination on glossy surfaces. Here, they compute and store SH coefficients of the spherical incident radiance distribution at cache record locations. For each such record, they also compute an approximate translational gradient, for each SH coefficient. Then, given shade point, they estimate the SH coefficients of incoming radiance using a first-order Taylor expansion, from which they can reconstruct the outgoing radiance towards camera. Although their approach successfully extends the irradiance caching paradigm to glossy surfaces, it is unable to properly treat glossy-glossy inter-reflection effects. The underlying dependence on the SH representation also suffers from ringing and aliasing in the reconstruction on highly glossy surfaces and an overhead when computing a large number of SH coefficients. Scherzer et al. [7] address the latter problem by accelerating radiance caching with a GPU-friendly representation. Here, they rely on cache records placed randomly on scene surfaces and use a fast pre-filtered MIPmap representation to store their radiance distribution, instead of a SH representation.

Our Contributions We devise a Hessian-based error metric suitable for radiance caching, and we apply a novel outgoing radiance-

*e-mail: yangyang.zhao2@mail.mcgill.ca
†e-mail: laurent@unity3d.com
‡e-mail: derek@cim.mcgill.ca

based approach to propose a new radiance caching method suitable for glossy-glossy inter-reflection effects, even in the presence of high-order glossiness. Specifically, we

- cache outgoing radiance, instead of the incident radiance distributions typically approximated with SH;
- propose a view-dependent record placement strategy that better matches the spatial variation of outgoing radiance;
- introduce generalize the Hessian-based error metric from irradiance caching [8] to both a view-dependent and view-independent placement strategy;
- derive new gradient formulations for outgoing and incident radiance that includes important terms missing from previous approximations; and,
- more accurately interpolate records placed on small curved objects, where we found SH representations to be beneficial.

3 Background

We briefly introduce the core concepts of caching-based methods (i.e., irradiance and radiance caching). The rendering equation [4] models the steady-state illumination at surface points x as

$$L(x,\omega_o) = L_e(x,\omega_o) + \int_\Omega L_i(x,\omega_i) f_r(\omega_i,\omega_o)(n\cdot\omega_i)\,d\omega_i\,, \qquad (1)$$

where ω_i and ω_o are incident and outgoing light directions at x, n is the surface normal, f_r is the BRDF, and $L_i(x,\omega_i)$ is the incident radiance distribution. The emitted light $L_e(x,\omega_o)$ is only non-zero for points x on light sources.

Indirect illumination refers to the contribution of $L_i(x,\omega_i)$ due to reflections from surfaces other than emitters. This contribution is difficult to compute due to the recursive nature of the integral. Irradiance and radiance caching accelerate this calculation by exploiting spatial correlation in the indirect illumination. The main idea here is to compute L_i accurately at only a sparse set of locations, from which we can interpolate approximate values of L_i at the remaining shade points in the scene.

Generally, these records are placed greedily: given a shade point x, we first search for records suitable for estimating L_i at x, according to a suitability-based distance criterion; if no suitable records are found, a new record is created at x and brute force Monte Carlo integration is used to compute the L_i value stored at this new record. A lazy-evaluation scheme referred to an *overture* pass is often employed to reduce duplicate lighting computations that arise from this greedy approach. The added complexity of an overture pass is negligible, and it does not impact the total number of created records.

Caching approaches generally reduce to solving three separate problems: *Cache Record Content*, *Cache Record Placement*, and *Shading*. *Cache Record Content* concerns the form of the information that is stored at each cache record, how the information is gathered and/or computed. *Cache Record Placement* concerns the strategy utilized when placing cache records. And *Shading* details the manner in which cache records are used to estimate the indirect illumination at a (non-cache) shading point.

3.1 Cache Record Content

Irradiance caching can only treat view-independent, diffuse reflection: no directional information is needed and we store the (scalar) irradiance. For radiance caching, we handle *view-dependent* glossy reflections, and so directional information (i.e., of the incoming or outgoing lighting) is required. Previous approaches typically store the incoming radiance distribution at cache records: Krivanek [5] uses SH to represent the incoming radiance distributions on surfaces, and Jarosz et al. [1] similarly use SH to store radiance in participating media. The SH representation has two main drawbacks: first, highly glossy inter-reflections require a large number of coefficients in order to accurately represent the potentially complex angular variations in the radiance; secondly, computing gradients of these SH radiance representations can be prohibitively inefficient. We instead choose to store outgoing radiance and detail the benefits of this design decision in the context of radiance caching.

3.2 Cache Record Placement

Most of the compute time in caching-based methods is spent on creating records. One typically aims to use as few records as possible while maintaining a desired image quality threshold. This is achieved by optimizing record placement and more accurately estimating the indirect illumination at shade points that fall outside these cache locations. As such, the manner in which we decide how to place cache records has a significant impact on the efficiency and accuracy of caching algorithms. Significant previous work has explored this problem, mostly in the context of irradiance caching.

Here, a fundamental concept associated with each cached record is that of its *valid region*. The valid region of a record refers to the region in space in which that record can influence shading points around it. The valid region of a record is usually defined as a (projected) disk or ellipse on surfaces around it.

Error-based Cache Record Placement Error-based approaches control record placement by thresholding on an upper bound of total error for points interpolated using the record inside its valid region. This error bound is defined by the integral of the difference between the correct illumination (which we never actually want to compute) and the illumination estimated by interpolation using the record. Compared to traditional heuristics-based placement strategies, such as the seminal *split-sphere* heuristic [10], error-based approaches can theoretically generate optimal[1] record placement. The various works in this area differ primarily in how they approximate the true error or, ultimately, how tight their upper bound on the total interpolation error can be.

In irradiance caching, the state-of-art [2,8] rely on a second-order Taylor expansion of the visibility-aware irradiance to approximate the actual indirect illumination. In radiance caching, Krivanek [5] provides a gradient formulation for the SH representation and experiments with a view-independent strategy that increases cache record density in regions of higher radiance gradients.

Motivated by the higher-order, visibility-aware variants used in irradiance caching, we devise a new error-based cache placement approach suitable for radiance caching, as well as a new view-dependent, Hessian-based outgoing radiance interpolation scheme.

4 View-dependent Radiance Caching Method

We begin with a review of technical details behind modern radiance caching methods, highlighting certain key limitations. We then present our reformulations and contrast them to alternative approaches, again in the context of the radiance caching problem.

Specifically, we present our view-dependent record placement strategies and discuss their advantages compared to traditional view-independent strategies (Section 4.1). We will detail the information stored at cache records (Section 4.2) followed by the gradient and Hessian computations needed to more accurately place and leverage these cache records (Section 4.3). Here, we will identify one key term (purposefully) omitted from previous gradient formulations, justifying a new practical approach to re-introduce it in the case of glossy-glossy interactions. In contrast to widely adopted spherical harmonics representations, we rely primarily on gradient and Hessian computations that can be performed semi-analytically. This approach drastically improves the performance of radiance caching

[1] In the sense of minimizing the interpolation error.

while reducing artifacts inherent to the SH-based approaches. Finally, we detail the technical aspects behnind realizing efficient computation of these first- and second-order derivatives, and their application during final shading (Section 4.4). Here, we leverage a hybrid record structure that can better handle shading variations on curved surfaces without incurring substantial cost.

4.1 Record Placement

We explore the placement of radiance cache records using a Hessian-based error metric. Similarly to prior work, we balance computation time and interpolation error by respecting an upper bound on the interpolation error over a cache record's valid region. The total error accrued by a cache record is defined as the integral of error due to extrapolating radiance over points in the record's valid region A. We control the size of a valid region implicitly by thresholding the error due to *all records* according to a user-defined parameter (the only parameter introduced by our method). Ours is a perceptually-motivated relative error metric, driven by studies that conclude that the human visual system is more sensitive to such errors.

Krivanek et al. [5] control the L_1 error of their SH approximation of the incoming radiance distribution at cache records, resulting in record placement that does not vary with the camera view. We, instead, control the error of our indirect illumination computation in a manner that depends on the reflection direction towards the camera (for every shading point). After exploring both view-independent and dependent alternatives, conclude that placing records according to the error of the view-evaluated outgoing radiance almost always makes more efficient and accurate use of the cached data.

4.1.1 View-dependent Caching Strategy

On glossy surfaces, the view-dependent outgoing radiance reflected towards the camera can be very sensitive to even slight variations in outgoing direction, particularly on highly glossy surfaces. Naturally, when rendering an image, we wish to minimize the error of only the portion of the indirect light field that is reflected towards the camera, and so we devise the following view-dependent error metric based on the view-evaluated outgoing radiance in order to make more effective use of our cached data:

$$\varepsilon^t = \iint_A \frac{|\, L(x_i+\Delta x) - L'(x_i+\Delta x)\,|}{L(x_i)} d\Delta x, \tag{2}$$

where $L(x_i+\Delta x)$ is the ground truth radiance at a point $x_i+\Delta x$ slightly offset from a cache point x_i, and $L'(x_i+\Delta x)$ is the approximate outgoing radiance at the offset point computed by extrapolation using cache records.

One key to improving placement here is to accurately estimate the variation of indirect illumination. We use a Hessian-based metric similar to that developed for irradiance caching [8], all while estimating final radiance using a second-order Taylor expansion about the cache record locations. We can estimate the total error ε^t using the 2×2 Hessian matrix of the outgoing radiance, parameterized on the tangent plane of the cache point. Specifically, we express the error using the two eigenvalues of the Hessian, u_1 and u_2, as:

$$\varepsilon^t \quad < \quad \tfrac{1}{2} \iint_A \left(\frac{|\, u_1 \,|}{L(x_i)} {x_e}^2 + \frac{|\, u_2 \,|}{L(x_i)} {y_e}^2 \right) dx_e dy_e, \tag{3}$$

where x_e and y_e are the projections of the offset vector Δx onto the first two eigenvectors of the Hessian. Given a user-defined maximum error threshold ε^{max}, we can devise anisotropic radii for cache records with elliptical valid regions similarly to [2], as:

$$\left(R_i^{u_1}, R_i^{u_2}\right) = \sqrt[4]{\frac{4\varepsilon^{max} L(x)}{\pi}} \left(\sqrt[4]{\frac{1}{|\, u_1 \,|}}, \sqrt[4]{\frac{1}{|\, u_2 \,|}} \right). \tag{4}$$

For such anisotropic records, we can directly compute the two axes of the elliptic valid region using the eigendecomposition of the outgoing radiance Hessian matrix.

4.1.2 View-independent Caching Strategy

We can generalize these metrics to apply to entire spherical radiance distribution, where cache record placement would now depend on the error estimation of indirect incident radiance at record locations. Since we will not consider any view-evaluated direction of the radiance distribution, such an approach would result in a view-independent error metric. The indirect incident radiance is a spherical distribution of radiance due to outgoing radiance from other surfaces in the scene. Here, as we are dealing with entire spherical distributions, we base our derivation on an SH representation. Note that, in practice, we very rarely fall-back to this view-independent scheme; the view-dependent cache is more generally effective in all but a few scenarios.

Our view-independent error is based on the L_1 error of the SH coefficient vector at a record. We can estimate the accumulated error inside a record's entire valid region as:

$$\varepsilon = \iint_A |\, \Lambda(x_i+\Delta x) - \Lambda'(x_i+\Delta x)\,|\; d\Delta x, \tag{5}$$

where $\Lambda(x_i+\Delta x)$ is the ground truth SH coefficients of the incident radiance and $\Lambda'(x_i+\Delta x)$ is our estimate of these coefficient obtained through interpolation. We once again apply a similar Hessian-based error metric, arriving at elliptical radii for our cache records using the Hessian of SH coefficients:

$$\left(R_i^{u_1}, R_i^{u_2}\right) = \sqrt[4]{\frac{4\,\varepsilon^{max}\,|\Lambda(x)|}{\pi}} \left(\sqrt[4]{\frac{1}{|\, u_1 \,|}}, \sqrt[4]{\frac{1}{|\, u_2 \,|}} \right), \tag{6}$$

where u_1 and u_2 are the two eigenvalues of H_x^{sum}, the sum of Hessian matrices accross each coefficient, i.e., $H_x^{sum} = \sum_{k=0}^{k=k_{max}} |H_x(\lambda_k)|$.

We can readily adapt this view-independent record placement scheme to treat *outgoing* radiance instead of *incident* radiance, simply substituting every instance of incident radiance above with outgoing radiance. This is possible since, given the incident radiance (whether ground truth or interpolated), one can obtain the outgoing radiance by spherical convolution with the BRDF at the shading point. As such, the only added complexity is in how we perform this transformation. We provide the details our application of a BRDF matrix for this task, later on.

We avoid this view-independent scheme except for select special cases (see below): while treating outgoing radiance instead of incident radiance (i.e., according to Krivanek's heuristic [5]) proves beneficial, it still completely suffers from significant drawbacks compared to the view-dependent approach we outline above. Namely, storing many SH coefficients per cache record, and computing the gradient and Hessian of these coefficients, incur large memory and computation overheads.

4.1.3 Analysis

We validate the view-dependent and view-independent outgoing radiance record placement strategies on a simple test scene comprising three intersecting glossy planes (i.e., forming half a box). The results in Figure 1 show that the view-independent strategy generates a nearly uniform distribution of records while our view-dependent strategy successfully adapts the record distribution according to the change in the indirect outgoing radiance.

4.2 Cache Record Content

When creating a new cache record, we store the record's position x_i in the scene and its surface normal n, and we compute and store estimates of the view-evaluated outgoing radiance $L(x_i) =$

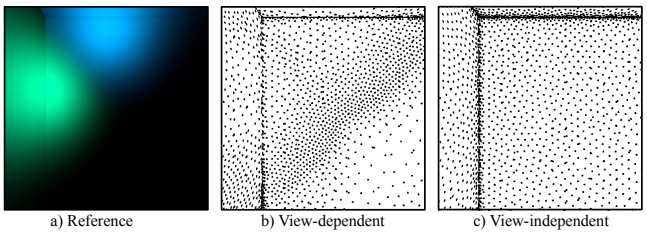

Figure 1: Comparison of view-dependent and view-independent strategies, visualizing indirect-only illumination at the corner of an entirely glossy box. The first image is the ground truth of indirect illumination, the middle illustrates records generated with our view-dependent strategy, and the right uses the view-independent outgoing radiance strategy. Both schemes use 2×10^3 records generated using a relative Hessian error metric.

$L(x_i, \omega)|_{\omega=\omega_o}$, the spatial gradient $\nabla_x L(x_i)$ and Hessian of outgoing radiance $H_x L(x_i)$, as well as the radius R of the record's valid region (or two radii R_1 and R_2 for anisotropic records).

In few situations (detailed in Section 4.3.4) we also rely on an SH encoding of the incoming radiance Λ_i, however we primarily interpolate outgoing radiance and only compute these SH coefficients lazily on narrow curved regions. In these few situations, we integrate the product of the incoming radiance and the BRDF to obtain the outgoing radiance.

We focus on interpolating indirect illumination on glossy surfaces and, in this scenario, we always employ a mixture of cosine-weighted and BRDF importance sampling to reduce the variance of any numerical integrals we compute using Monte Carlo.

When numerically projecting incident radiance distribution onto the SH basis functions, to compute its SH coefficients, we compute solve for each projection coefficients integral

$$\lambda_l^m = \int_\Omega L_i(x, \omega_i)\, y_l^m(\omega_i)\, d(\omega_i)\,, \tag{7}$$

using Monte Carlo integration with uniform hemispherical samples, where λ_l^m is the m^{th} band-l spherical harmonic coefficient and $y_l^m(\omega_i)$ is the associated spherical harmonics basis function.

The outgoing radiance distribution L can also be approximated with an SH vector Λ_o using the incident radiance's projection coefficient vector Λ_i, using the approach outlined by Pharr and Humphreys [6]: given the $(l_{\max}+1)^2 \times (l_{\max}+1)^2$ *BRDF SH product matrix* $\mathbf{M}$, the outgoing radiance SH vector is simply $\Lambda_o = \mathbf{M}\Lambda_i$. Since the matrix M depends only on the BRDF, we can precompute M for every BRDF in our scene and store these matrices for later use. We can easily evaluate the outgoing radiance towards, e.g., the viewer in direction ω_o given its projection coefficient vector Λ_o as $L_o(x, \omega)|_{\omega=\omega_o} = \mathbf{M}\Lambda_i Y(\omega_o)$.

As discussed earlier, SH provide a smooth reconstruction and natural support for surface orientation changes, however accurately reconstructing high-frequency incoming radiance variations requires costly higher-order SH expansions. This not only increases the number of coefficients and the size of $\mathbf{M}$, but since each coefficient and its derivatives are computed separately, the overall computation and storage overhead is further impacted. For these reasons, we fall-back to an SH representation, albeit a novel one that relies on view-evaluated outgoing radiance instead of incident radiance, sparingly. Figure 2 compares renderings of the same scene with highly glossy surfaces using our analytic method and an SH-based representation. Note that the latter fails due to severe ringing, even using an order ten SH expansion. As such, we rely primarily on directly storing and interpolating outgoing radiance instead of reconstructing it using SH. The remainder of our exposition will focus on deriving the mathematical models for the outgoing radiance method, including computing the necessary gradients and Hessians for it.

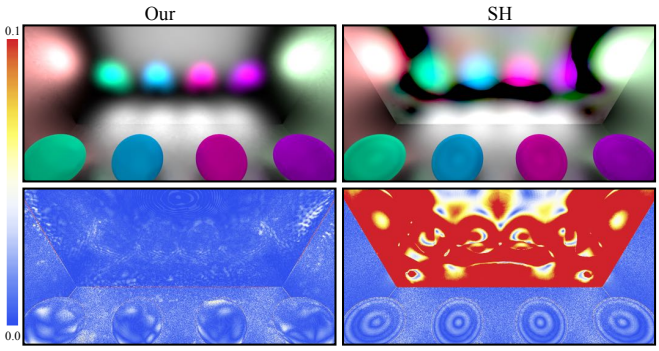

Figure 2: We tested our SH approach and direct analytic outgoing radiance approach on a scene with glossy Phong surfaces with glossiness exponents of 100, and both images only illustrate the indirect illumination. Left to right: the result rendered using the analytic outgoing radiance method, followed by the result using an SH representation of outgoing radiance. Both images use the same number of cache records (roughly 8K). We can observe visible ringing artifacts in the SH approach.

4.3 Gradient and Hessian

Radiance caching requires spatial gradients and Hessians of the outgoing radiance at the record locations, and these are used for record placement and interpolation during shading. Accurately estimating the translation derivatives of outgoing radiance (and of its SH projection coefficients) is crucial to the quality of the final images (i.e., using a fixed number of cache records). Our approach focuses only on the translation derivatives since the rotational gradient (and Hessian) is prohibitively costly to compute. This simplification is common to even the irradiance caching scenario [2]. We introduce two methods to compute the derivatives of outgoing radiance, depending on the underlying representation: an analytic approach and one that relies on a SH representation. Recall that, in most situations, we utilize the analytic approach as it is more efficient and accurate than the SH approach.

4.3.1 Glossy-Glossy Gradient

Before introducing our methods, let us focus on the term $L_i(x, \omega_i)$ in Equations 1 and 7. Previous work [5] assume that the incident radiance arriving from other surfaces towards a record is spatially invariant, cancelling the term chain term $\nabla_x L_i(x, \omega_i)$ from the computation of the entire gradient. Although this assumption seems to work well in scenes where $L_i(x, \omega_i)$ is due solely to diffuse (inter-)reflectors, it no longer holds in the case of glossy-glossy inter-reflections, e.g., the glossy Cornell box in Figure 3.

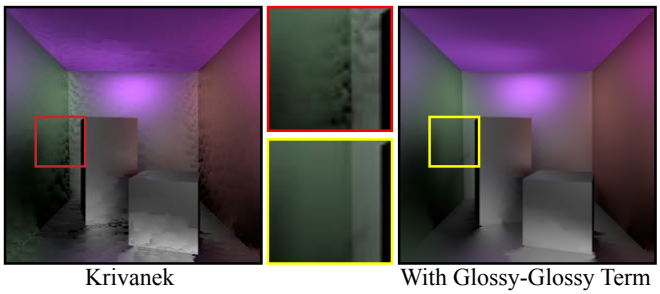

Figure 3: The rendering on the left uses standard radiance caching [5]. Here, we observe significant artifacts, whereas the image on the right incorporates the missing glossy-glossy gradient chain term. Both images are rendered with approximately 4K cache records, and the new gradient formulation requires approximately 10% additional rendering time (376s compare to 333s).

We name this missing gradient term the *Glossy-Glossy term*. To better demonstrate the impact of this term on the entire radiance caching pipeline, we consider a simple illustrative test scene in Figure 4: the consists of only two surfaces and a single point light source. The bottom surface observed by the camera is a glossy surface, and the vertical surface (not directly visible from the camera) can be either diffuse or glossy. The contribution from the missing spatial gradient term can be non-negligible when dealing with glossy-glossy reflections and, in order to correctly sample and reconstruct glossy-glossy reflections, the gradient/Hessian derivations and computations from previous work must be augmented with a more accurate estimate of the spatial derivative of $L_i(x, \omega_i)$.

4.3.2 Analytic Translation Derivative

We convert from the solid angle formulation of Equation 1 to a surface-area parameterization to facilitate the derivative derivations, below:

$$L(x, \omega_o) = \int_S L(y \to x)\, f_r(x, \omega_i, \omega_o)\, G(x \leftrightarrow y)\, dy \tag{8}$$

where y is the position of a surface visible from x in incident direction ω_i, $L_i(y \to x)$ is the reflected radiance from y to x, and:

$$G(x \leftrightarrow y) = \frac{-(n \cdot \mathbf{r})(n_y \cdot \mathbf{r})}{|\,\mathbf{r}\,|^4}, \mathbf{r} = x \to y \tag{9}$$

is the geometry Jacobian term that results from the change of parameterization.

Then, the spatial gradient $\nabla_x L(x, \omega_o)$ and the Hessian $H_x L(x, \omega_o)$ are:

$$\int_S \nabla_x (\, L_i(y \to x)\, f_r(x, \omega_i, \omega_o)\, G(x \to y)\,)\, dy, \tag{10}$$

and

$$\int_S H_x (\, L_i(y \to x)\, f_r(x, \omega_i, \omega_o)\, G(x \to y)\,)\, dy. \tag{11}$$

The derivative inside the integral can be decomposed into several terms using the chain rule, and we need to compute the first- and second-order derivatives of each of these three terms: the spatial derivative of the geometry term $G(x \to y)$, which is straightforward to derive from prior work; the spatial derivative of the BRDF term, for which there are no analytic forms for an arbitrary BRDF f_r; and the spatial derivatives of the incident radiance L_i, which can also be challenging to solve for. We discuss these two remaining cases, below.

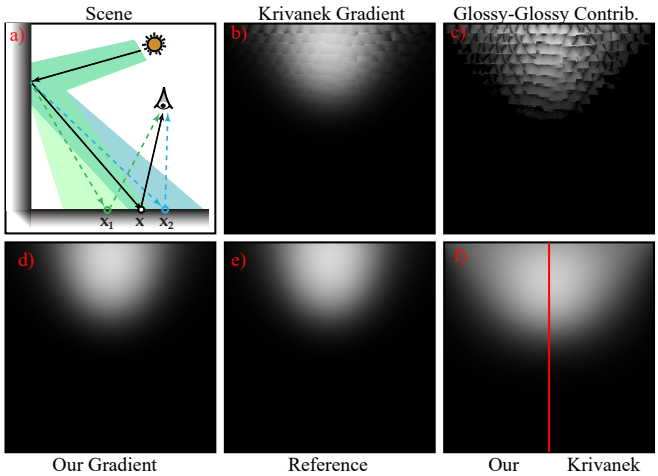

Figure 4: a) Example scene setup: we first render the scene with traditional radiance caching [5]. The result shown in b) has severe artifacts. In c) we show the contribution from the glossy-glossy gradient. In d) we render this scene with our gradient formulation and compare with the reference image in e). After changing the vertical surface to be diffuse, we render the scene with traditional radiance caching and compare to our approach. The results in f) show that these two approaches produce similar results when glossy-glossy gradient contributions are negligible.

In the case of the spatial gradients and Hessians of the BRDF, if we restrict ourselves to commonly used reflection models with analytic forms (such as the Lambertian, Phong and Microfacet models), we can derive the gradient and Hessian in closed form. As the camera location and hemispherical sample y are fixed, the two angles ω_i and ω_o in the BRDF term $f_r(\omega_i, \omega_o)$ are only dependent on the position of the record x. We parameterize the BRDF function with the record location x and compute its derivative analytically using a symbolic mathematics package, for the three analytic BRDF reflections models mentioned above.

The last subterm we need to consider is the gradient and Hessian of the incident radiance $L_i(y \to x)$ reflected from y to x. If y is located on a diffuse surface, L_i will be (approximately) spatially constant and its derivatives will be negligible. However, this no longer holds for glossy surfaces and assuming so in radiance caching can cause significant visual artifacts. To derive the gradient and Hessian of the glossy-glossy incident radiance term, we first express L_i as an integral of an additional (i.e., third) bounce of light (i.e., from the camera) as:

$$L_i(y \to x) = \int_\Omega L_i(y, \omega_i^y)\, f_r(y, \omega_o^y, \omega_i^y)\, (n_y \cdot \omega_i^y)\, d\omega_i^y. \tag{12}$$

Above, $L_i(y, \omega_i)$ is the incident radiance from direction ω_i at y. We can safely assume that ω_i remains constant due to spatial changes at y since the location of the hemispherical sample does not change, and so the spatial derivative of $L_i(y, \omega_i)\,(n \cdot \omega_i)$ is zero. Then, we need only compute the spatial gradient and Hessian of the BRDF $f_r(y, \omega_o^y, \omega_i^y)$ of this tertiary bounce.

To do so, we approach the problem similarly as we did for the spatial BRDF gradient from the secondary bounce, above: we first reparameterize the BRDF function in Cartesian coordinates, then arrive at an analytic formulation of $\nabla_x (f_r(\omega_o, \omega_i))$ and $H_x (f_r(\omega_o, \omega_i))$ that can be solved using symbolic computation (again, assuming analytic reflection models at the secondary and tertiary bounces).

For Microfacet reflection models, we further assumethat the Fresnel term and shadowing factors remain spatially constant across the cache record's valid region. Then, we need only derive the spatial gradient and Hessian for the normal distribution function. We provide automatically code-generated output from a symbolic computation package for the widely adopted GGX microfacet distribution as well as the Phong reflection model.

4.3.3 SH Translational Derivative

Our analytic approach avoids computing derivatives for arbitrary BRDFs by restricting the type reflection models (and, so, surfaces) we support. We can also avoid computing derivatives of arbitrary BRDFs by relying on a precomputed SH intermediate representation. Here, we need to compute the spatial derivative of $\mathbf{M}\Lambda_i Y(\omega_o)$. If the BRDF function does not vary across the surface, the the BRDF matrix remains constant and we only need the derivative of SH basis functions and the SH coefficients of the incident radiance. For the SH basis functions, we derive spatial gradients by reparameterization from the solid angle to the surface-based forms. Then, we can compute translational gradients and translational Hessians based on this parameterization. The spatial derivative of incident radiance are obtained relatively easily: we rewrite Equation 7 as a surface integral and derive the gradient and Hessian directly from it. The derivatives of the glossy-glossy term can be computed using the same approach outlined above, now without limiting the surface reflection model.

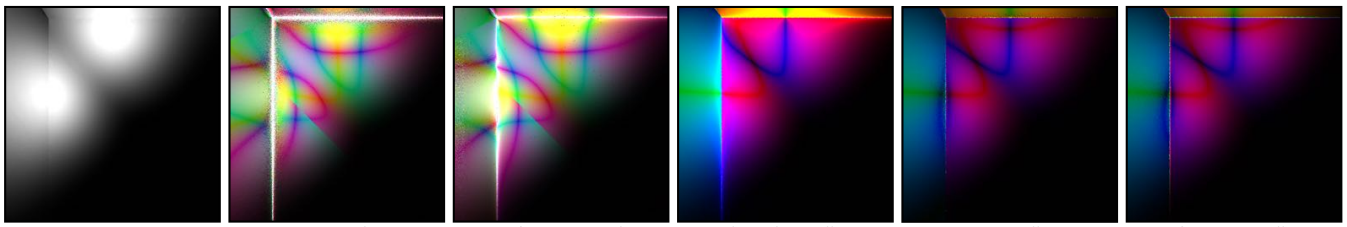

Figure 5: We visualize the gradient and Hessian of the radiance bounced toward camera. We compared the color mapped result from our approach and previous approach. Then we compare them with the reference gradient and Hessian computed with finite difference method. We visualize the gradient in world coordinate with the mapping $(dL/dx, dL/dy, dL/dz) \rightarrow (r,g,b)$. Hessian matrices are mapped in following way: $(d^2L/dx^2, d^2L/dy^2, d^2L/dxdy) \rightarrow (r,g,b)$.

4.3.4 Rotational Derivative

To accurately interpolate the radiance on curved surfaces we need to estimate radiance changes due to by normal variation. In theory, it is possible to accurately estimate such rotational derivatives for this purpose, however, we encountered some important difficulties here: first, the rotational derivatives for glossy BRDFs can become intractable to solve in closed form (even with advanced symbolic computation packages). For simple reflection models where we obtain closed-form solutions, the resulting formulations are computationally inefficient to evaluate and fail to offer an accurate enough estimation when the BRDF is too high frequency. In order to ignore the rotational gradients, our radiance interpolants can only operate in regions with tight maximum normal deviation. On very curved surfaces, where this assumption is broken, the SH approach offers a smoother filtered reconstruction, and so we rely on interpolating the outgoing SH vector using its gradient information. As such, we employ a hybrid caching structure to deal with curved surfaces, where we compute SH coefficients for the records located on curved surfaces and, during rendering, we interpolate these coefficients (instead of the outgoing radiance) at shade points. Since the number of SH records needed are usually very small, this has little impact on the overall rendering time.

4.3.5 Analysis

Figure 5 shows that our methods yields relatively accurate approximations of the first- and second-order derivatives of outgoing radiance. Our SH approach requires computing the derivatives for each coefficient. With the number of coefficients sufficient to approximate the radiance distribution, the compute time per cache record is usually much longer than in the analytic case.

4.4 Shading

Given a shade point x, we can compute outgoing radiance by interpolating the nearby cached records. To improve this estimation, we used a second-order Taylor expansion about the nearest cache points to more accurately estimate the coefficients at point x:

$$L_o \approx \frac{\sum_S w_i(x)(L_i + \nabla_x L_i \Delta x_i + \frac{1}{2}\Delta x_i^T H_x(L_i)\Delta x_i))}{\sum w_i}, \qquad (13)$$

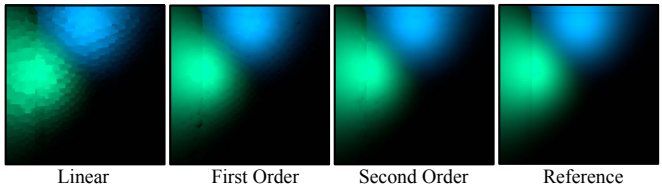

Figure 6: Comparison of different orders of Taylor extrapolation on a test scene. Using the same number of cache records (0.7K), second-order extrapolation clearly outperforms other approaches.

where S is the set of all records with x in their valid regions and with normal deviation less than θ_{max}, and w_i is the weight of the record which depends on the distance and normal deviation from point x, as:

$$w_i(x) = \max(0, 1 - k_t * k_t) * \max(0, k_r)^{16}, \qquad (14)$$

where 16 is an empirical value that generates smooth interpolation on curved surfaces in practice, and k_t and k_r are translation and rotation weights:

$$\begin{aligned} k_t &= \sqrt{\left[\frac{\Delta x_i \cdot v^{u1}}{R_i^{u1}}\right]^2 + \left[\frac{\Delta x_i \cdot v^{u2}}{R_i^{u2}}\right]^2 + \left[\frac{\Delta x_i \cdot n_i}{R_i^{u1}}\right]^2} \\ k_r &= n_i \cdot n_x . \end{aligned} \qquad (15)$$

In Figure 6, second-order extrapolation shows clear improvement compared to first-order extrapolation, but the improvement brought by second-order extrapolation is less prominent when the scene has many small surfaces.

5 Results and Implementation

We implemented our new radiance caching algorithm on top of PBRT [6] framework version 2.0. The tests are performed on a PC with Intel i7 3770K CPU running Windows 7. All results are rendered with a sampler render from PBRT with 16 samples per pixel. The maximum allowed normal deviation is 0.1 rad. The lower and upper bounds of record radius are respectively set to the projection of one pixel and 10% of image size.

Figure 7 is an equal-time comparison of our analytical Hessian-based error metric against the spherical harmonic Hessian-based error metric. Each image is rendered at a resolution of 512×512, and only one bounce indirect illumination is shown. For each cache

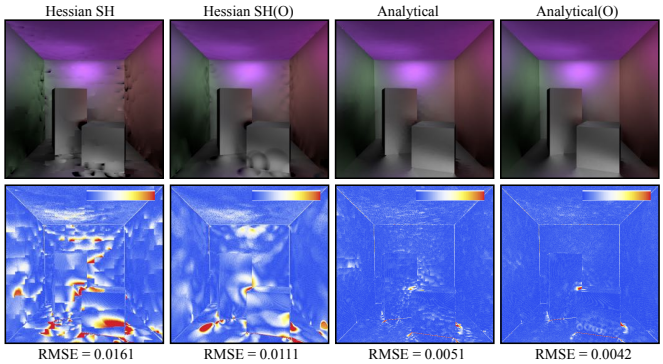

Figure 7: We show an equal-time comparison (200s) of our analytical approach and the spherical harmonic approach with and without an additional overture pass. We also show the result of path tracing with the same rendering time as reference. While the overture pass offers 30% error reduction for the spherical harmonic approach, it is still far away from the quality of our analytical approach.

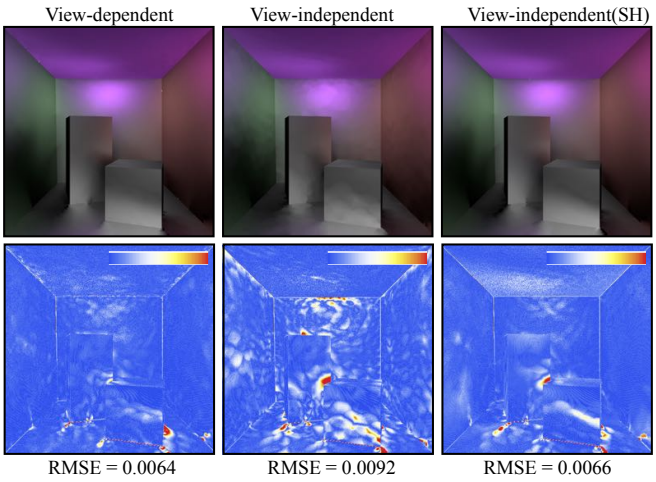

Figure 8: We compare the view-dependent and view-independent record placement strategy with the same amount of cache record used. The first image is the result of our view-dependent strategy with 2K records. With the outgoing radiance interpolation, it clearly outperforms the view-independent record placement strategy with the same amount of records. The right most image is the result of the view-independent strategy with spherical harmonic interpolation using 2K cache records.

record, we use 8192 hemispherical sampling rays. We adjust the maximum allowed error for radiance caching to achieve similar rendering times. With 200s rendering time, the spherical harmonic approach is only able to generate 0.8K cache records, and shows severe artifacts. Our analytical approach has much better results with 5K cache records generated within the same amount of time. We also improve two approaches with an overture pass. We control the parameters to ensure both approaches spend approximately 200s to generate cache records. Because of the spherical harmonic evaluation, the rendering pass of spherical harmonic approach is around 20s longer than analytical approach. Our analytical approach is still twice better than the spherical harmonic approach even if the overture pass brings significant improvements to the spherical harmonic approach. We also notice that the overture pass has little impact on our approach compared to SH approach.

To evaluate the performance of our view-dependent record placement strategy, Figure 8 shows the results of these two strategies with $2K \pm 2\%$ records. Similar to the equal-time comparison above, the resolution of image is 512×512, only one bounce indirect illumination is rendered. The new view-independent strategy results in much better images with the same amount of records with outgoing radiance interpolation. Also the view-dependent record placement strategy with radiance interpolation has slightly better quality than the

Figure 9: Two scenes with GGX microfacet surfaces rendered by our radiance caching. The left side image is the the modified Kitchen scene. The right side is the Cornell box scene and corresponded records distribution.

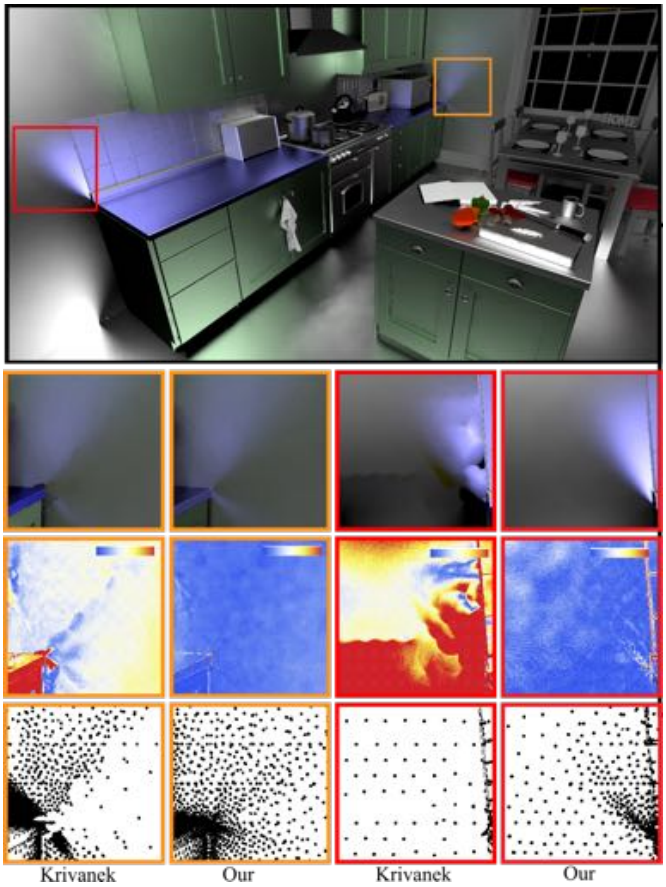

Figure 10: We show the advantage of our new method on a relatively complex kitchen scene in dealing both with sharp indirect reflections and highly glossy surfaces. Our approach captures and adjusts the record placement properly to produce high-quality results. The details from comparison shows that the view-independent spherical harmonic approach suffers severe artifacts and systematical errors on high order glossy surfaces.

view-independent strategy with SH coefficient interpolation despite the fact the view-dependent records are much slower to generate.

Figure 9 shows our radiance caching works with GGX microfacet reflection models. Compare to the phong reflection model, microfacet models produce more sharp reflections. With a little bit more records, our view-dependent strategy successfully adapt and capture these indirect shadows without additional clamping and excessive placement of records.

Figure 10 demonstrates the results of our approach on a kitchen scene with highly glossy surfaces. This scene has some relatively complex glossy-glossy reflections between high-order Phong surfaces (up to order 400). The resulting sharp indirect shadows pose serious challenges for previous radiance caching methods. We compare our approach against other approaches on two close-ups of this scene. The resolution of a full image is 1024×768, and only the one-bounce indirect illumination is rendered. For each cache record, we use 14400 hemispherical gather rays. 90K records are generated for the full image. The total rendering time is 100 minutes. The close-ups are rendered with 1.6K records for every approach. The comparison shows clearly that our approach can not only accurately capture the sharp indirect reflections, but also avoid the inaccurate radiance reconstructions caused by spherical harmonics.

In Figure 11, we focus on the details of two metal pots in the kitchen scene. We compare the straightforward outgoing radiance interpolation method with a hybrid interpolation scheme. The com-

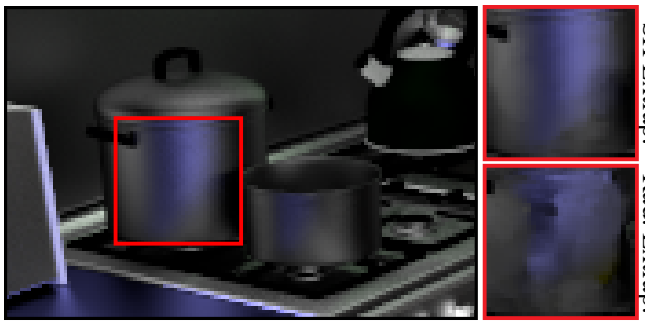

Figure 11: The detail of metal pots in the kitchen scene. The spherical harmonic approach is able to provide much better interpolation on curved surfaces.

parison on the right side of the figure shows that interpolating the harmonic (SH) coefficients on curved surfaces produces better details on curved parts of the scene.

Figure 12 shows a scene with the Austrian Imperial crown with a highly glossy surface (order 1000 Phong). The reflection of the crown on the glossy surface is very bright and sharp. The images are rendered at 1280×768 pixels. The reference is generated with path tracing with 32768 samples per pixel. Our radiance caching algorithm uses approximately $100K$ records, each record is generated by 32768 gather rays. The rendering time of our radiance caching is 2.63 hours compared to 30 hours using path tracing.

6 Conclusion and Future Work

In this paper, we propose a new radiance caching algorithm by combining and extending previous work on irradiance caching and radiance caching. We introduce and discuss a new Hessian-based view-dependent record placement strategy in a radiance caching context. We interpolate the outgoing radiance and avoid the spherical hamonics approximation along with the drawbacks it causes. Our results show that storing radiance improves both the quality and computation time of radiance caching. We discuss a problem overlooked during derivative computation in previous work. We correct this oversight and propose our new analytical derivative formulations for radiance caching. Our new approach extends radiance caching to highly glossy surfaces. We proved that our approach is able to detect the variance of radiance bounced towards camera and place records accordingly.

Despite the significant improvements brought by our new radiance caching approach. There are still a few unanswered questions. Although the spherical harmonics approach has been proven inefficient when rendering single image, it is unknown whether this view-independent approach can be efficient to render animations of a static scene. The outgoing radiance towards a new camera location must be recomputed instead of reconstructed from SH coefficients. Applying a texture variation is trivial since it can be multiplied after interpolating outgoing radiance. Our method depends on the parameters of BRDF remaining locally unchanged. As a result, we cannot deal with non-constant BRDFs. Our hybrid cache offers a simple, low cost fix to interpolation artifacts on curved surfaces. The lighting information from the occluded region of the hemisphere after normal variation is also unaccounted for. A practical and accurate rotational derivative would greatly improve our method here.

Another fundamental problem that exists in all current radiance caching and irradiance caching methods is that they use a greedy algorithm to place cache records. The records in highly occluded shadow areas with large radius will potentially cause artifacts outside the shadow. For example in Figure 7, the minor artifact near the shadow of small box is caused by this problem. To solve this problem, we need a new record placement method for both algorithms.

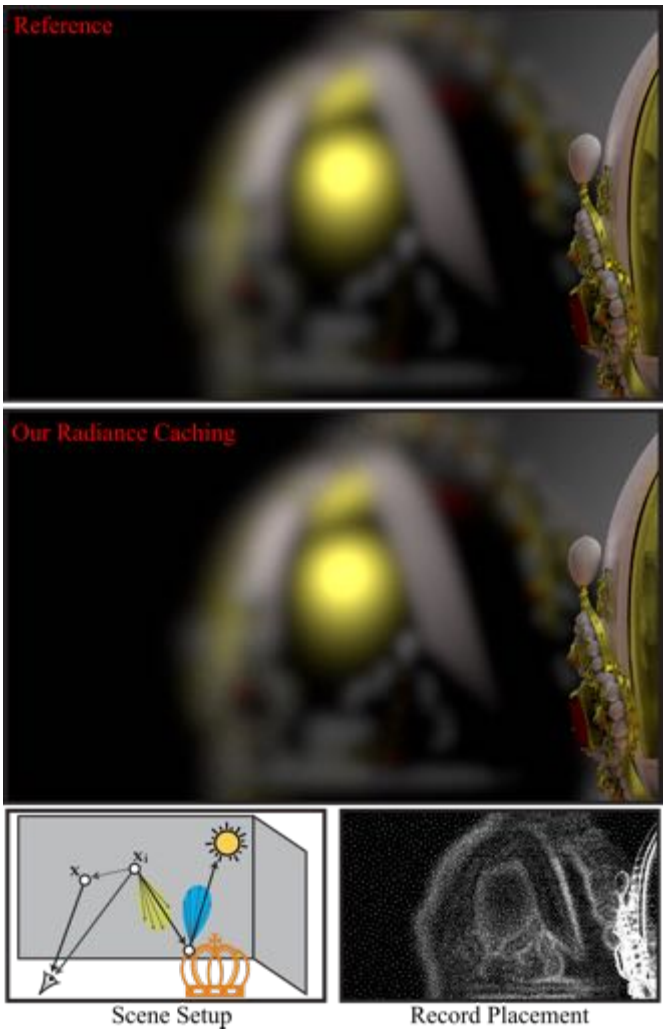

Figure 12: The reflection of the Austrian Imperial Crown on a highly glossy surface. The top left image is the illustration of the setup of the scene. The bottom left image shows the record placement that our approach generated. Our radiance caching is able to capture the details on the reflection and produces high-quality results. Model courtesy by Martin Lubich.

References

[1] W. Jarosz, C. Donner, M. Zwicker, and H. W. Jensen. Radiance caching for participating media. *ACM Transactions on Graphics (Presented at SIGGRAPH)*, 27(1):7:1–7:11, Mar. 2008. doi: 10.1145/1330511.1330518

[2] W. Jarosz, V. Schonefeld, L. Kobbelt, and H. W. Jensen. Theory Analysis and Applications of 2D Global Illumination. *ACM Transactions on Graphics*, 31(5):243–253, 2012.

[3] H. W. Jensen. Global Illumination using Photon Maps. In *Rendering Techniques '96*, pp. 21–30. Springer-Verlag, 1996.

[4] J. T. Kajiya. The Rendering Equation. *SIGGRAPH Comput. Graph.*, 20(4):143–150, Aug. 1986.

[5] J. Krivanek. *Radiance Caching for Global Illumination Computation on Glossy Surfaces*. Ph.d. thesis, Université de Rennes 1 and Czech Technical University in Prague, December 2005.

[6] M. Pharr and G. Humphreys. *Physically Based Rendering, Second Edition: From Theory To Implementation*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2010.

[7] D. Scherzer, C. H. Nguyen, T. Ritschel, and H.-P. Seidel. Pre-convolved Radiance Caching. *Computer Graphics Forum (Proc. EGSR 2012)*, 4(31), 2012.

[8] J. Schwarzhaupt, H. W. Jensen, and W. Jarosz. Practical Hessian-Based

Error Control for Irradiance Caching. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH Asia 2012)*, 31(6):193–193, 2012.

[9] G. Ward and P. Heckbert. Irradiance Gradients. In *Eurographics Rendering Workshop*, pp. 85–98, May 1992.

[10] G. J. Ward, F. M. Rubinstein, and R. D. Clear. A Ray Tracing Solution for Diffuse Interreflection. *SIGGRAPH Comput. Graph.*, 22(4):85–92, 1988.